
Rule WLM358: QUEUING IN IOS WAS A MAJOR CAUSE OF I/O DELAY

Finding: Queuing in the I/O Supervisor (IOSQ) was a major cause of the I/O delay with the volume.

This finding applies only to MVS versions prior to OS/390 Release 3, and to MVS versions with OS/390 Release 3 if I/O Priority Management has **not** been specified.

Impact: This finding may have a MEDIUM IMPACT or HIGH IMPACT on the performance of the device.

Logic flow: The following rules cause this rule to be invoked:

- Rule WLM350: I/O activity may have caused significant delays
- Rule WLM351: I/O activity may have caused significant delays
- Rule WLM352: I/O activity may have caused significant delays for server service class
- Rule WLM353: I/O activity may have caused significant delays for server service class

Discussion: IOSQ time is the time from the issuance of a STARTIO macro until the Start SubChannel (SSCH) instruction is issued. After the STARTIO macro is issued, the software determines whether the device is busy with the system on which the STARTIO macro was issued (that is, whether there is an available Unit Control Block (UCB) for the device). If the device is not busy with this system, the SSCH instruction is issued. However, if the device is busy with this system (a UCB is available), the I/O request is queued. Thus, IOSQ time always means that the device is unable to handle additional requests from this system.

This discussion of IOSQ time does not always apply to Parallel Access Volumes (PAVs)¹. With PAV devices, MVS creates multiple UCBs for each device, depending on how many “alias devices” have been defined. The multiple UCBs allow multiple active concurrent I/Os on a given device when the I/O requests originate from the same system². Using PAVs can dramatically improve I/O performance by nearly eliminating IOSQ.

¹PAV devices are available with Enterprise Storage Server (ESS). With PAV devices, a “base device” address is defined, and a UCB is associated with this base address. “Alias device” addresses can be defined and UCBs are associated with the alias device addresses.

²Multiple Allegiance allows multiple active concurrent I/O operations on a given device when the I/O requests originate from different systems. The Multiple Allegiance feature is available with Enterprise Storage Server (ESS).

The following example illustrates the output from Rule WLM358:

RULE WLM358: DEVICE IOS QUEUING TIME WAS A MAJOR CAUSE OF DASD DELAYS

A major part of the potential I/O delay to the TSO Service Class could be attributed to queuing in the I/O Supervisor (IOSQ). IOSQ time is caused by too many I/O operations directed to the device or lengthy device response times (perhaps caused by high seeking, by high RPS delays, or by high PEND time. Please refer to the WLM Component User Manual for advice on how to minimize device IOSQ time.

Suggestion: Large IOSQ times usually involve the following situations:

- Multiple data sets may be active on the volume. This situation is the most common and easiest to solve. The data sets can be redistributed among different logical volumes, to eliminate the queuing for the single volume.
- The data sets can be placed on PAV devices or redistributed among different logical volumes, to eliminate the queuing for the single volume.
- If using static PAVs, assign more aliases to the device.
- If using dynamic PAV, increase the number of PAVs associated in the pool for the subsystem.
- Ensure that all PAVs that should be bound to the device are online and are operational. You can use the DEVSERV QP and DS QP,xxxx,UNBOX commands to do this.
- Multiple users may be using the same data set on the volume. Depending upon the data set characteristics, duplicate copies of the data set placed on different volumes may solve the IOSQ problems.
- Multiple application systems may be using the volume experiencing high IOSQ times. In this case, perhaps application redesign or scheduling can solve the problem.
- A particular application (or system function) may be executing I/O to the device faster than the device can respond.
- The overall device response time (PEND, DISC, and CONN) times may be large, such that the device is unable to provide quick response to the I/O requests. This situation will be revealed by large values in the PEND, DISC, or CONN measures. Consider moving files to a faster storage (coupling facility structure, expanded storage, Data In Memory, etc.).

Also, consider speeding up or reducing the I/O on the path or the device
(e.g., specify optimal VSAM options, revise blocking options, etc.).